arXiv:1602.09061v1 [cs.LO] 29 Feb 2016

# Bounded Model Checking of Pointer Programs Revisited*

Witold Charatonik[1] and Piotr Witkowski[2]

[1] Institute of Computer Science, University of Wroclaw,
wch@cs.uni.wroc.pl
[2] pwit@pwit.info

**Abstract.** Bounded model checking of pointer programs is a debugging technique for programs that manipulate dynamically allocated pointer structures on the heap. It is based on the following four observations. First, error conditions like dereference of a dangling pointer, are expressible in a fragment of first-order logic with two-variables. Second, the fragment is closed under weakest preconditions wrt. finite paths. Third, data structures like trees, lists etc. are expressible by inductive predicates defined in a fragment of Datalog. Finally, the combination of the two fragments of the two-variable logic and Datalog is decidable.
In this paper we improve this technique by extending the expressivity of the underlying logics. In a sequence of examples we demonstrate that the new logic is capable of modeling more sophisticated data structures with more complex dependencies on heaps and more complex analyses.

## 1 Introduction

Automated verification of programs manipulating dynamically allocated pointer structures is a challenging and important problem.

In [11] the authors proposed a bounded model checking (BMC) procedure for imperative programs that manipulate dynamically allocated pointer structures on the heap. Although in this procedure an explicit bound is assumed on the length of the program execution, the size of the initial data structure is not bounded. Therefore, such programs form infinite and infinitely branching transition systems. The procedure is based on the following four observations. First, error conditions like dereference of a dangling pointer, are expressible in a fragment of first-order logic with two-variables. Second, the fragment is closed under weakest preconditions wrt. finite paths. Third, data structures like trees, lists (singly or doubly linked, or even circular) are expressible in a fragment of monadic Datalog. Finally, the combination of the two-variable fragment with Datalog is decidable. The bounded model checking problem for pointer programs is then reduced to the satisfiability of formulas in the combined logic. The authors gave an algorithm solving the satisfiability problem in 2NExpTime.

In this paper we further develop this method. We formulate a general BMC problem (which was not formulated earlier; it was only applied in a rather ad hoc

* Research supported by NCN Grant no. 2011/03/B/ST6/00346

manner) and show that BMC of null pointers from [11] is its instance. We also provide several other instances (Examples 5–9). The logic used in [11] was quite restrictive, in particular it allowed no existential quantifiers and no sharing; it allowed only one Datalog program; it could not express cardinality constraints; it had very limited support for content analysis (only a single query to the Datalog program was possible); it was unable to speak about inductive (i.e., defined with Datalog) properties in postconditions. Here, by extending the expressivity of the underlying logics we are able to model more sophisticated data structures with more complex dependencies on heaps and more complex analyses. Specifically, in a series of examples we show the change in expressivity due to new features of our logic: use of existential quantifiers (and other relaxations to the syntax, like multiple queries to Datalog programs) in Examples 3, 4 and 8; sharing structure in Examples 2, 3 and 8; use of multiple Datalog programs in Examples 8 and 9; cardinality constraints in Example 3; support for content analysis in Example 4; inductive properties in both pre- and post-states in Example 8.

It is worth noting that a bounded pointer program can traverse only a bounded fragment of a data structure, which suggests that there is no point in allowing data structures of unbounded size. However, describing only traversable fragments is not enough for some analyses that check certain properties of the whole heap. In Examples 3 (where we analyze cardinality constraints on the heap as a whole) and 9 (where we detect memory leaks) the result of the analysis depends on parts of the heap that are not touched by the program. These are properties of the heap as a whole, and not of the traversable fragment.

## 2 Related work

There are many approaches to automated verification of pointer programs that emerged recently. Most of them use logical formalisms to describe heaps, capture program semantics and verify (partial) program correctness by Hoare method. It requires expressing combinations of heap shape with data properties, and quickly leads to undecidable logics. Powerful proof engines (employing abstractions, theorem proving and/or SMT based reasoning) [47,39,29] are then used to find proofs for specification obtained in that way allowing even the *full functional verification* [46,30,38,35] of data structures, the holy grail of software analysis. Our approach is different. We aim at bounded model checking, which allows to find bugs rather than to prove their absence, and the logic we use is decidable, what certainly limits its expressivity. A common belief is that pure first-order logic is too weak to reason even about the simplest data structures. However the result from [13] implies that the two-variable logic with counting $C^2$, a decidable first-order fragment, is in some cases sufficient. There a combination of $C^2$ with Datalog was defined and shown to be reducible simply to $C^2$. Here we continue this line of research by demonstrating expressivity of another combination of $C^2$ with Datalog that translates directly to a decidable extension of $C^2$ called $C^2$ with trees [14,43]. Our longer term aim is to continue investigation on expres-

sive yet decidable extensions of first order fragments applied to pointer program verification. Some specific related work is discussed below.

*Abstract interpretation and shape analysis.* One possibility is to compute over-approximations of the set of reachable states and to represent them, together with program actions, as formulas of a 3-valued logic with transitive closure. This is the approach taken e. g., in [42,22,45,31,2]. Its soundness relies on abstract interpretation, but it may result in false positives, i. e., an erroneous state that is unreachable from any of starting state may belong to the over approximation of the set of reachable states. On the contrary, logics with Datalog are expressive enough to precisely model reachable states of simple pointer programs. However, these programs may as well be abstractions of other pointer programs and abstract interpretation techniques might be applicable in our setting. A notable difference between our approach and the one mentioned above is that a model of a formula with Datalog represents concrete state of a program, while a model of 3-valued formula represents (an abstraction of) a set of states. Recent work on 3-valued abstractions aims at verifying heap shape and data stored there [18,17]. For simplicity, in our approach only heap shape is represented, but since we represent heaps as models of formulas that admit unary predicates, these can be used to simulate finite domain data.

*The Pointer Assertion Logic Engine.* Another option is to use monadic second order logic on trees [25,32]. Sets of states of pointer programs are modeled using graph types, which consist of tree backbones with some additional edges. As observed in [11], structures defined by Datalog programs can be seen as tree backbones, but in our approach additional edges may be specified in a fragment of first order logic, while graph types specify these additional edges in a dynamic logic. The employed logic is powerful, but of non-elementary complexity. In contrast, our decidable logics with Datalog are relatively weak, but of an elementary, NExpTime complexity. This means that not all graph types are expressible in our logics. On the other hand, due to arbitrary binary predicates and to presence of cardinality constraints, our logics are not subsumed by MSOL on trees.

*Separation logic.* A powerful, but undecidable formalism for local reasoning about pointer programs with lists was introduced in [34]. This was a kind of proof-theoretical approach to program verification, since proofs of Hoare triples must be manually constructed using an intuitionistic proof calculus. On the contrary, both approaches discussed above as well as ours rely on decidable logics. Decidable fragments of separation logics are also studied [4,9,6,7,16,3], including fragments with not only lists but also general inductively defined predicates [33,23,1,36,8]. Complexity of these fragments vary from NP through ExpSpace up to nonelementariness. Our logic, $C^2$ + Datalog includes a semantic restriction bsr that, roughly, forces separation of data structures defined by a single Datalog program. This suggests a relation between our approach and the above mentioned. However, logics with Datalog may also express structures that

intersect, provided that they are defined by different Datalog programs in a formula.

*Logic for reachable patterns in linked data structures.* It seems that in terms of expressibility the logic most related to ours is the one from [44]. It admits arbitrary Boolean combinations of reachability constraints similar to universally quantified guarded formulas. The exact difference in expressive power needs to be investigated, but the two logics differ in terms of complexity and underlying decision procedures. The satisfiability problem for the logic in [44] has NExpTime lower bound and elementary upper bound and it is also based on a translation to a kind of monadic second-order logic on trees (the authors say that they have another doubly-exponential procedure, but it is not published). Although an arbitrary number of universally quantified variables is allowed in the logic, the formulas in two examples provided in [44] use at most three variables. Moreover, three variables are used only to define properties that constitute semantic restrictions in our logic. We are able specify all shapes occurring there in our combination of $C^2$ with Datalog. Moreover, although admitting only two variables, our logic allows us unguarded quantification and counting quantifiers.

A significant difference between our approach to BMC of pointer programs and many other BMC techniques [5,24,15,21] is that we bound only the length of program paths to be symbolically executed and not the size of the input data structures. Thus we are able to perform model checking for infinite-state transition systems.

In addition to combined verification of heap shape with data [18,17,36], heap shape with size [33] and balanceness of data structures [19] one may verify heap shape together with content properties, where these properties are specified as description logic formulas or UML diagrams [10,26]. As we show in Section 3.3 our logic allows us to specify size constraints as well as content properties, but not balanceness. Other graph logics embedable in $C^2$ for modeling heap shape as well as content properties were earlier considered [27,28,40]. These logics are incomparable with ours, since they aim at representing abstractions of program states (a model represents a set of heaps), while logics with Datalog are designed to represent concrete states (a model represents a heap).

## 3 Two-variable logic with counting and Datalog

In this section we introduce our logic. It is a two-variable fragment of the first-order logic extended with counting quantifiers and inductive predicates in form of Datalog programs. The obtained logic is decidable by reduction to two-variable logic with counting and trees [14] and is expressive enough to model interesting properties of dynamically allocated pointer structures.

### 3.1 Monadic Datalog Programs

Datalog is a declarative logic programming language. Syntactically it is a subset of Prolog that does not use function symbols of arity greater than 0 (i. e., con-

stants). It is often used as a query language for deductive databases. Here we use it to extend the expressive power of some logics to be able to define dynamic structures on heap.

Let $\Sigma_E$ and $\Sigma_I$ be disjoint signatures, the former (called *extensional signature*) containing relational symbols of arity at most 2, equality and constants, and the latter (*intensional signature*) containing only unary symbols. Signature $\Sigma_I$ defines symbols that occur in heads of clauses from a Datalog program, while remaining symbols occurring in clauses come from $\Sigma_E$. We will call them *intensional* (respectively *extensional*) symbols. Following [11], we are interested in *monadic Datalog programs*. A clause in such a program is a Horn clause where the only positive literal has a unary predicate in its head and there are additional constraints on remaining literals, as stated below.

**Definition 1.** *A* monadic Datalog program *over $\Sigma_E$ and $\Sigma_I$ is a finite set of clauses of the form*

$$p(u) \leftarrow B(u) \wedge \bigwedge_{i=1}^{l} [r_i(u, v_i) \wedge q_i(v_i)], \text{ where}$$

1. *$p(\cdot), q_1(\cdot), \ldots, q_l(\cdot)$ are $\Sigma_I$-predicates;*
2. *$r_1(\cdot,\cdot), \ldots, r_l(\cdot,\cdot)$ are distinct $\Sigma_E$-predicates;*
3. *$B(u)$, is a (possibly empty) quantifier-free first-order $\Sigma_E$-formula containing only constants and the variable $u$;*
4. *$l \geq 0$ and $u, v_1, \ldots, v_l$ are distinct variables.*

*Monadic Datalog programs are further called Datalog programs for short. Datalog programs will be denoted by blackboard bold letters $\mathbb{P}, \mathbb{Q}, \mathbb{R}$, $\mathbb{P}_{\text{list}}$ etc.*

Consider an example Datalog program $\mathbb{P}_{\text{list}}$ from Figure 1. The extensional signature of $\mathbb{P}_{\text{list}}$ is $\{\text{next}(\cdot,\cdot), =, \text{NULL}\}$ and intensional signature is $\{\text{list}(\cdot)\}$. Our intention is that $\text{list}(x)$ denotes that $x$ is a node of a singly linked list, where every node is either `NULL` or has one successor pointed to by next pointer.

$$\mathbb{P}_{\text{list}} = \{\ \text{list}(x) \leftarrow \text{next}(x, y) \wedge \text{list}(y),\quad \text{list}(x) \leftarrow x = \text{NULL}\ \}.$$

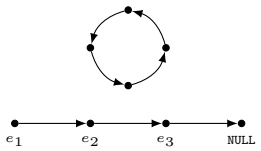


Fig. 1: Datalog program $\mathbb{P}_{\text{list}}$ and a structure $\mathcal{M}$. Edges represent the relation $\text{next}(\cdot,\cdot)$.

Datalog programs have natural least fixed point semantics. Given a relational structure $\mathcal{M}$ over $\Sigma_E$ and a Datalog program $\mathbb{P}$ over $\Sigma_E$ and $\Sigma_I$ the *least extension of $\mathcal{M}$ w. r. t. $\mathbb{P}$* is the least $\Sigma_E \cup \Sigma_I$ structure $\mathcal{M}_{\mathbb{P}}$ such that 1) $\mathcal{M}$ is

contained in $\mathcal{M}_{\mathbb{P}}$, and 2) if $[p(u) \leftarrow B(u) \wedge \bigwedge_{i=1}^{l}[r_i\,(u, v_i) \wedge q_i(v_i)] \in \mathbb{P}$ and $\mathcal{M}_{\mathbb{P}} \models B(e) \bigwedge_{i=1}^{l}[r_i\,(e, e_i) \wedge q_i(e_i)]$ then $\mathcal{M}_{\mathbb{P}} \models p(e)$, for all $e, e_1, \ldots, e_k \in \mathcal{M}$. Consider the Datalog program $\mathbb{P}_{\text{list}}$ and structure $\mathcal{M}$ from Figure 1. The least extension of $\mathcal{M}$ w. r. t. $\mathbb{P}_{\text{list}}$ is the structure $\mathcal{M}_{\mathbb{P}_{\text{list}}} = \mathcal{M} \cup \{\text{list}(\texttt{NULL}), \text{list}(e_3), \text{list}(e_2), \text{list}(e_1)\}$. The nodes in the cycle, are not members of the list; although the structure $\mathcal{M} \cup \{\text{list}(e) \mid e \in \mathcal{M}\}$ satisfies conditions 1) and 2) above, it is not the least one.

For a given Datalog program $\mathbb{P}$ let $\Sigma(\mathbb{P})$ be the subset of $\Sigma_E$ containing all binary predicates mentioned in point 2 of Definition 1. Let $\mathbb{P}_1$, $\mathbb{P}_2$ be Datalog programs over $\Sigma_E$ and $\Sigma_I^1$, $\Sigma_E$ and $\Sigma_I^2$ respectively. Programs $\mathbb{P}_1$, $\mathbb{P}_2$ are called *disjoint* if $\Sigma_I^1 \cap \Sigma_I^2 = \emptyset$. Let $\mathbb{P} = \mathbb{P}_1, \ldots, \mathbb{P}_k$ be a sequence of pairwise disjoint Datalog programs such that, for every $i \in 1, \ldots, k$, the extensional vocabulary of $\mathbb{P}_i$ is $\Sigma_E$. The semantics of Datalog program sequence $\mathbb{P}$ is the same as semantics of a Datalog program $\bigcup_{i=1}^{k} \mathbb{P}_i$: for a structure $\mathcal{M}$ over $\Sigma_E$ we define its least extension w. r. t. sequence $\mathbb{P}$ as $\mathcal{M}_{\mathbb{P}_1,\ldots,\mathbb{P}_k} = \mathcal{M}_{\bigcup_{i=1}^{k} \mathbb{P}_i}$.

### 3.2 Syntax and semantics of the logic

The logic we use in bounded model checking of pointer programs combines the two variable logic with counting and inductive predicates defined by Datalog programs. The two variable logic with counting ($\text{C}^2$) is a decidable fragment of first order logic containing formulas whose all subformulas have at most two free variables, but may contain counting quantifiers of the form $\exists^{\geq k}, \exists^{=k}, \exists^{\leq k}$. With these quantifiers one may specify that there are at least, precisely or at most $k$ elements with a given property.

We employ $\text{C}^2$ formulas over vocabulary $\Sigma_E \cup \Sigma_I$, but we impose some restrictions on $\Sigma_I$ atoms that occur in these formulas. Let $\phi$ be a $C^2$ formula over $\Sigma_E \cup \Sigma_I$ and let $\phi'$ be its negational normal form. We say that a $\Sigma_I$-atom $p(x)$ has a restricted occurrence in $\phi$ if either $p(x)$ occurs positively in $\phi'$ and only in in scope of existential quantifiers or $p(x)$ occurs negatively in $\phi'$ and only in scope of universal quantifiers. For example $p(x)$ has a restricted occurrence in formulas $\forall x\ p(x) \rightarrow \psi$, $\forall x\ (p(x) \wedge q(x)) \rightarrow \psi$, $\exists_x\ p(x) \wedge \psi$ or $\exists_x\ p(x) \wedge q(x) \wedge \psi$, where $\psi$ is some $C^2$ formula with one free variable $x$ and no occurrence of $p(x)$, and $q(\cdot)$ is some $\Sigma_I$-predicate. An occurrence of atom $p(x)$ in formula $\forall_y \exists_x p(x) \wedge \psi$ is not restricted, because $p(x)$ occurs positively and in scope of a $\forall$ quantifier.

**Definition 2 (Syntax of $\text{C}^2$+Datalog).** *An expression $[\mathbb{P}_1, \ldots, \mathbb{P}_k]\phi$ is a* $\text{C}^2$+Datalog *formula over $\Sigma_E$ and $\Sigma_I$ if*

1. *$\mathbb{P}_1, \ldots, \mathbb{P}_k$ are pairwise disjoint Datalog programs,*
2. *the extensional (respectively intensional) vocabulary of $\mathbb{P}_i$ is contained in $\Sigma_E$ (respectively in $\Sigma_I$), for $i \in \{1, \ldots, k\}$,*
3. *$\phi$ is a formula of the two-variable logic with counting quantifiers over the signature $\Sigma_E \cup \Sigma_I$, and*
4. *every $\Sigma_I$-literal occurring in $\phi$ is an intensional literal defined by $\mathbb{P}_1$ or $\mathbb{P}_2$, or is a constant literal or has only restricted occurrences in $\phi$.*

Notice that Datalog programs $\mathbb{P}_1$ and $\mathbb{P}_2$ are *privileged*, i. e., $(\Sigma_I^1 \cup \Sigma_I^2)$-predicates may form arbitrary constant- or non-constant literals in $\phi$. On the contrary, literals made of predicates defined by remaining Datalog programs may either be constant literals or have only restricted occurrences in $\phi$.

From now on, when a vocabulary $\Sigma_E$ is clear from context, we will write $\mathrm{const}(v)$ as a shortcut for the formula $\bigvee_{c \in \Sigma_E} v = c$. Let $\mathbb{P} = \mathbb{P}_1, \ldots, \mathbb{P}_k$ be a sequence of disjoint Datalog programs such that $\mathbb{P}_i$ is over $\Sigma_E$ and $\Sigma_I^i$, for $i \in \{1, \ldots, k\}$. For a given $\Sigma_E$-structure $\mathcal{M}$ let $\mathcal{M}_\mathbb{P}$ be the least extension of $\mathcal{M}$ w. r. t. sequence $\mathbb{P}$. We say that $\mathcal{M}_\mathbb{P}$ obeys the *bounded-sharing* restriction (bsr for short) if $\mathcal{M}_\mathbb{P}$ is a model of all sentences of the form

$$\forall u_1, u_2, v \ \ (s_1(u_1, v) \wedge s_1(u_2, v) \wedge u_1 \neq u_2 \rightarrow \mathrm{const}(v)) \ \text{ and}$$

$$\forall u_1, u_2, v \, (s_1(u_1, v) \wedge s_2(u_2, v) \rightarrow \mathrm{const}(v)) \,,$$

where $s_1$ and $s_2$ are two distinct predicates occurring in $\Sigma(\mathbb{P}_i)$, for $i \in \{1, \ldots, k\}$. We say that $\mathcal{M}_\mathbb{P}$ obeys the *bounded intersection* restriction (bir for short) if for all distinct predicates $p(\cdot), q(\cdot) \in \Sigma_I^i$, structure $\mathcal{M}_\mathbb{P}$ models

$$\forall u. p(u) \wedge q(u) \rightarrow \mathrm{const}(u).$$

Intuitively, the bounded-sharing restriction says that two pointers occurring in the same Datalog program cannot point to the same memory cell. The restriction ensures that data structures defined by a single Datalog program are tree-like, in the sense that in-degree of their nodes is $\leq 1$. Additionally, the bounded intersection restriction forces these structures to be disjoint. In both cases an exception is made for constant nodes; they can model e. g., the `NULL` node which is unique and shared among all data structures on heap, or the first common node in two lists that have a common suffix.

**Definition 3 (Semantics of $\mathrm{C}^2$ + Datalog).** *Let $[\mathbb{P}]\phi$ be a formula of $\mathrm{C}^2$ + Datalog over $\Sigma_E$ and $\Sigma_I$, where $\mathbb{P} = \mathbb{P}_1, \ldots, \mathbb{P}_k$ and let $\mathcal{M}_\mathbb{P}$ be a finite structure over $\Sigma_E \cup \Sigma_I$ such that*

- *$\mathcal{M}_\mathbb{P}$ is the least extension of some $\Sigma_E$-structure $\mathcal{M}$ w. r. t. Datalog program sequence $\mathbb{P}$,*
- *$\mathcal{M}_\mathbb{P}$ satisfies bounded-sharing and bounded-intersection restrictions, and*
- *$\mathcal{M}_\mathbb{P} \models \phi$.*

*Then $\mathcal{M}_\mathbb{P}$ is said to satisfy $[\mathbb{P}]\phi$, in symbols $\mathcal{M}_\mathbb{P} \models [\mathbb{P}]\phi$.*

Although both bounded-sharing and bounded-intersection restrictions are expressible in our logics they cannot be removed, as they are crucial in the satisfiability decision procedure in [43]. With these restrictions we may express many data structures including lists and trees, also with limited sharing of substructures (see examples in the next section), but we cannot express arbitrary DAGs.

Sometimes we would like to define linear constraints on the number of realizations of unary predicates in a structure $\mathcal{A}$. When the vocabulary $\Sigma$ is known

from the context and $p_1(\cdot), \ldots, p_l(\cdot)$ are unary symbols from $\Sigma$ we write $\Delta$ to denote a system of linear (in)equalities in variables $\#p_1, \ldots, \#p_k$. We say that $\mathcal{A}$ satisfies $\Delta$ (written $\mathcal{A} \models \Delta$) if the valuation $\rho$ defined as $\rho(\#p_i) = |p_i^{\mathcal{A}}|$ satisfies $\Delta$. Here $|p_i^{\mathcal{A}}|$ denotes the number of elements of structure $\mathcal{A}$ that satisfy the predicate $p_i$. Let $[\mathbb{P}_1, \ldots, \mathbb{P}_k]\phi$ be a $\mathrm{C}^2 + \mathrm{Datalog}$ formula and $\Delta$ be a system of linear (in)equalities over intensional predicates from $\mathbb{P}_1$ and $\mathbb{P}_2$, and over unary extensional predicates from $\Sigma_E$. We write $[\mathbb{P}_1, \ldots, \mathbb{P}_k, \Delta]\phi$ for a formula with the same semantics as the starting $\mathrm{C}^2 + \mathrm{Datalog}$ formula, but with the additional requirement that $\mathcal{M}_{\mathbb{P}_1,\ldots,\mathbb{P}_k} \models \Delta$.

The following theorem was proven in [43], where $\mathrm{C}^2 + \mathrm{Datalog}$ was called $\mathrm{C}^2_{\mathrm{r2}} + \mathrm{Datalog} + \{\mathbf{bsr}, \mathbf{bir}\}$.

**Theorem 1 ([43], Cor. 3.27).** *Finite satisfiability problem for* $\mathrm{C}^2 + \mathrm{Datalog}$*, even enriched with linear (in)equalities, is* NExpTime*-complete.*

The requirement that we allow at most 2 privileged Datalog programs in $\mathrm{C}^2 +$ Datalog formulas cannot be easily removed. It is related to an open problem, whether satisfiability for the two-variable logic $\mathrm{FO}^2$ with more than two successors of finite linear orders is decidable (note that we may express two successors of finite linear orders in $\mathrm{C}^2 + \mathrm{Datalog}$).

### 3.3 Modeling data structures in $\mathrm{C}^2 + \mathrm{Datalog}$

Let us demonstrate expressive power of $\mathrm{C}^2 + \mathrm{Datalog}$. We do it by writing a handful of formulas describing heaps of imperative pointer programs. Here we show examples of data structures, and in Section 4.3 we give examples of analyses that can be modeled. More examples can be found in the PhD thesis of the second author [43].

A heap can be seen as a relational structure, where nodes are heap elements (we assume that all elements of the heap are of the same size), binary predicates denote pointers between nodes, constants denote nodes pointed to by program variables and there is a distinguished constant `NULL` denoting the null value. Binary relations are interpreted as partial functions (*functionality restriction*) — although every pointer on a heap has some value, in our setting we allow it to have no value at all. Moreover, we sometimes introduce auxiliary unary and binary predicates to express additional properties. Note that the property of a binary predicate $f(\cdot,\cdot)$ being a partial function is easily expressible in our logic by a formula $\forall x \exists^{\leq 1} y f(x, y)$.

The logic $\mathrm{C}^2$+Datalog on structures that satisfy bounded-sharing and bounded-intersection restrictions strictly subsumes the logics considered in [11,13]. Therefore, after recalling the simplest examples from [11], we present exemplary structures not expressible in the subsumed logics. Many other examples can be found in [11,13,43]. Let us start with a simple example of a singly-linked list.

*Example 1.* The simplest linked data structure is a singly-linked `NULL`-terminated list with head in some specified node h. For $\Sigma_E = \{\mathrm{next}(\cdot,\cdot), \mathrm{h}, \mathtt{NULL}\}$ and

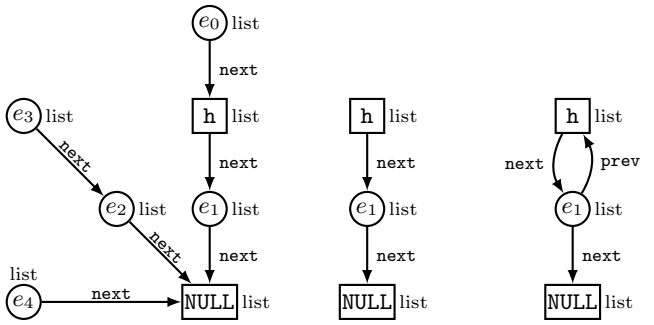


Fig. 2: Models of $[\mathbb{P}_{list}]\phi$ from Example 1 for different formulas $\phi$. Here $\phi_1 = (\forall u\, \neg\text{next}(u,h) \wedge \forall u \forall v\, \text{next}(u,\texttt{NULL}) \wedge \text{next}(v,\texttt{NULL}) \rightarrow u \approx v)$ and $\phi_2 = \forall u \forall v\, (u \neq \texttt{NULL} \wedge v \neq \texttt{NULL}) \rightarrow (\text{next}(u,v) \leftrightarrow \text{prev}(v,u))$.

$\Sigma_I = \{\text{list}(\cdot)\}$ let $\varphi = [\mathbb{P}_{\text{list}}]\phi$ be a $\text{C}^2$ + Datalog formula where $\mathbb{P}_{\text{list}}$ is defined in Figure 1. By defining $\phi$ to be just a query list(h) we force models of $\varphi$ to contain a list from $h$ to NULL made of next edges. One of the possible models of $[\mathbb{P}_{\text{list}}]\phi$ is depicted in Figure 2a. Thanks to bounded-sharing restriction NULL and h can be the only nodes shared by different lists — in Fig. 2a node NULL is shared while h is not. Moreover, the functionality restriction ensures that every node emits at most one $\text{next}(\cdot,\cdot)$ pointer. The bounded-intersection plays no role here, since there is only one intensional predicate in the signature.

We may ensure that $h$ is indeed the head of the list (and not an internal node) by adding a conjunct $\forall u\, \neg\text{next}(u,h)$ to $\phi$. Moreover, to ensure that the list with head in $h$ is the only list in the structure, we add to $\phi$ formula $\exists^{\leq 1} u\, \text{next}(u,\texttt{NULL})$. This is depicted in Figure 2b. We can further modify our formula to capture doubly linked lists by adding to $\phi$ a conjunct $\forall u \forall v\, (u \neq \texttt{NULL} \wedge v \neq \texttt{NULL}) \rightarrow (\text{next}(u,v) \leftrightarrow \text{prev}(v,u))$ as in Figure 2c.

The NULL node represents the undefined memory address. It may be pointed to by an arbitrary number of pointers, but no pointer can start in it. This is expressed by a formula $\forall u\, \neg\text{next}(\texttt{NULL},u)$ in the context of the example above, and in general by $\bigwedge_{r(\cdot,\cdot)\in\Sigma_E} \forall u\, \neg r(\texttt{NULL},u)$. We assume that such a conjunct is implicitly included in every formula we write here.

The next example shows the difference between modeling a data structure with a single Datalog program and a sequence of Datalog programs.

*Example 2.* Consider the following Datalog programs.

$$\mathbb{P}_{llist} = \{\ \text{llist}(x) \leftarrow \text{left}(x,y) \wedge \text{llist}(y),\ \text{llist}(x) \leftarrow x = \text{NULL}\ \}.$$
$$\mathbb{P}_{rlist} = \{\ \text{rlist}(x) \leftarrow \text{right}(x,y) \wedge \text{rlist}(y),\ \text{rlist}(x) \leftarrow x = \text{NULL}\ \}.$$

The formula $[\mathbb{P}_{llist}, \mathbb{P}_{rlist}](\text{llist}(\text{h}_1) \wedge \text{rlist}(\text{h}_2))$ expresses structures where $\text{h}_1$ is a node on a NULL-terminated list made of $\text{left}(\cdot,\cdot)$ pointers and $\text{h}_2$ is a node on a

NULL-terminated list made of right$(\cdot,\cdot)$ pointers. These two lists may be disjoint, like in Fig. 3 (left), or may share nodes, even non-constant ones, like in Fig. 3 (right). There may also be other lists in the structure. Notice the difference between $[\mathbb{P}_{llist}, \mathbb{P}_{rlist}](\text{llist}(\text{h}_1) \wedge \text{rlist}(\text{h}_2))$ and $[\mathbb{P}_{llist} \cup \mathbb{P}_{rlist}](\text{llist}(\text{h}_1) \wedge \text{rlist}(\text{h}_2))$. The latter one forbids sharing of non-constant nodes, and therefore the structure in Fig. 3 (left) is one of its models while the one in Fig. 3 (right) is not.

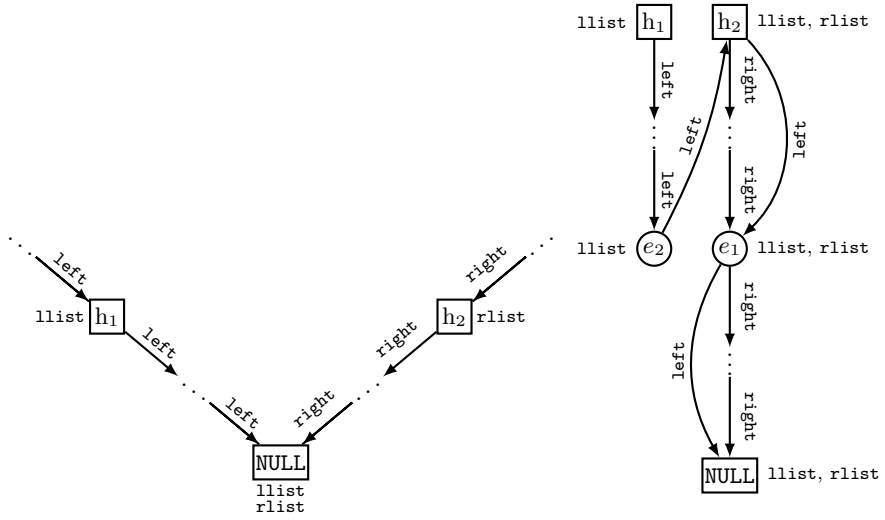


Fig. 3: Models of a formula $[\mathbb{P}_{llist}, \mathbb{P}_{rlist}](\text{llist}(\text{h}_1) \wedge \text{rlist}(\text{h}_2))$ from Ex. 2. Lists on the left structure are disjoint, with the exception of NULL node. Lists on the right structure share non-constant node $e_1$ and constant $\text{h}_2$. Dots denote arbitrary number of intermediate nodes.

The formulas written so far did not employ global cardinality constraints, but, since all of them consist of at most two Datalog programs, they may be supplemented by such. Consider the example below.

*Example 3.* Let us define heaps being two binary trees rooted in $\text{r}_1$ and $\text{r}_2$ respectively. We require that the number of nodes shared by these two trees is the half of their size. A C$^2$+Datalog formula encoding the property is $[\mathbb{TREE}_1, \mathbb{TREE}_2, \Delta]\phi$, where

$$\begin{aligned}\mathbb{TREE}_1 = \{\ & \text{tree}_1(u) \leftarrow \text{left}_1(u,v) \wedge \text{tree}_1(v) \wedge \text{right}_1(u,w) \wedge \text{tree}_1(w),\\ & \text{tree}_1(u) \leftarrow u \approx \text{NULL}\ \},\end{aligned}$$

and $\mathbb{TREE}_2$ is just $\mathbb{TREE}_1$ with every subscript 1 replaced by 2. The formula $\phi$ is the conjunction of

$$\begin{gathered}\mathrm{tree}_1(\mathrm{r}_1) \wedge \mathrm{tree}_2(\mathrm{r}_2)\\ \forall u.\ \mathrm{tree}_1(u) \rightarrow (u \approx \mathrm{r}_1 \vee \exists v.\ ((\mathrm{left}_1\ (v,u) \vee \mathrm{right}_1\ (v,u)) \wedge \mathrm{tree}_1\ (v)))\\ \forall u.\ \mathrm{tree}_2(u) \rightarrow (u \approx \mathrm{r}_2 \vee \exists v.\ ((\mathrm{left}_2\ (v,u) \vee \mathrm{right}_2\ (v,u)) \wedge \mathrm{tree}_2\ (v)))\\ \forall u.\mathrm{shared}(u) \leftrightarrow \mathrm{tree}_1(u) \wedge \mathrm{tree}_2(u).\end{gathered}$$

The global cardinality constraint $\Delta$ is just a single equation $\{\#_{\mathrm{tree}_1} + \#_{\mathrm{tree}_2} = 2 * \#_{\mathrm{shared}}\}$.

Existence of both trees is guaranteed by $\mathbb{TREE}_1$, $\mathbb{TREE}_2$ and the first conjunct of $\phi$. Next two conjuncts express that $\mathrm{r}_1$ ($\mathrm{r}_2$) must be reached from every node labeled by $\mathrm{tree}_1(\cdot)$ (respectively $\mathrm{tree}_2(\cdot)$). This effects in that all nodes labeled by $\mathrm{tree}_1(\cdot)$ (respectively $\mathrm{tree}_2(\cdot)$) belong to the tree rooted in $\mathrm{r}_1$ ($\mathrm{r}_2$). The last conjunct of $\phi$ defines auxiliary predicate $\mathrm{shared}(\cdot)$ to label exactly the nodes shared by both trees. Then, the required cardinality constraint is expressed by $\Delta$.

In the examples above we analyzed only shape or quantitative properties of heaps. A novel approach to verification of pointer programs was recently proposed in [10], where some properties of heap content are formally specified as description logic formulas or UML diagrams, and heap shape is defined in a fragment of separation logic. The last example in this section shows that both content and shape properties may be expressed in $\mathrm{C}^2 + \mathrm{Datalog}$.

*Example 4 (Information system of a company, the running example from [10]).* A software company is divided into departments, has a number of employees (some of them are managers), who work for projects (some of which are large) and projects are ordered by clients. There are certain number restrictions on relations between these entities, as specified by UML diagram in Fig. 4, e. g., each employee works for at most one project, while each project has an arbitrary number of employees working on it. The diagram also establishes a subsumption relation between large projects and projects (i. e., every large project is a project) and similarly for managers and employees. Projects, employees, departments and clients are stored on `NULL`-terminated lists on $\mathrm{next}(\cdot,\cdot)$ pointers. The information system of the company manipulates these lists; it may add and remove their nodes, assign managers to departments and projects etc. Every such an operation must preserve properties expressed by UML diagram. In this example we focus only on defining in $\mathrm{C}^2 + \mathrm{Datalog}$ a heap shape and a part of the diagram concerning projects, managers and employees. We write a formula $[\mathbb{P}_{list}]\phi$. First, we have a standard list definition, the Datalog program $\mathbb{P}_{list}$ as in Example 1. The formula $\phi$ expresses that

1. constant nodes pHd and eHd are heads of two lists; $\mathrm{list}(\mathrm{pHd}) \wedge \mathrm{list}(\mathrm{eHd}) \wedge \forall u.\ (\neg\mathrm{next}(u, \mathrm{pHd}) \wedge \neg\mathrm{next}(u, \mathrm{eHd}))$,
2. projects and employees are stored in some lists on heap; $\forall u.\mathrm{project}(u) \vee \mathrm{employee}(u) \rightarrow \mathrm{list}(u)$,

3. all nodes on the list headed in pHd (respectively eHd) are projects(respectively employees); project(pHd)$\wedge\forall u.\,(\text{project}(u) \rightarrow u \approx \texttt{NULL} \vee \exists v.\,(\text{next}(u,v) \wedge \text{project}(v)))$ and a similar formula for employee$(\cdot)$,
4. all projects (respectively employees) are on list headed in pHd (respectively eHd);
   $\forall u.\,(\text{project}(u) \rightarrow u \approx \text{pHd} \vee \exists v.\,(\text{next}(v,u) \wedge \text{project}(v)))$ and a similar formula for employee$(\cdot)$,
5. projects and employees are disjoint; $\forall u.\neg\text{project}(u) \vee \neg\text{employee}(u)$
6. each employee has at most one pointer worksFor$(\cdot,\cdot)$ to a project, indicating a project that the employee is working on (recall that we are writing formula of a logic with functionality restriction); $\forall u\forall v.\text{employee}(u) \wedge \text{worksFor}(u,v) \rightarrow \text{project}(v)$,
7. employees have a Boolean field is_manager$(\cdot)$ marking them as managers; $\forall u.\text{is_manager}(u) \rightarrow \text{employee}(u)$,
8. similarly, projects have a Boolean field is_large$(\cdot)$ marking them as large projects; $\forall u.\text{is_large}(u) \rightarrow \text{project}(u)$,
9. each project has at most one pointer managedBy$(\cdot,\cdot)$ to an employee being its manager; $\forall u\forall v.\text{project}(u) \wedge \text{managedBy}(u,v) \rightarrow \text{is_manager}(v)$,

Conjuncts 1—5 express heap shape properties, i. e., that the heap consists of two disjoint lists of projects and employees, while conjuncts 6—9 define properties of heap content, i. e., a fragment of the UML diagram. One can also include in $\phi$ other properties of the information system, not expressed by the UML diagram but encodable in the logic, like

10. the manager of a project works for the project;
    $\forall u\forall v.\text{project}(u) \wedge \text{managedBy}(u,v) \rightarrow \text{worksFor}(u,v)$,
11. at least 10 employees work on each large project
    $\forall u.\text{is_large}(u) \rightarrow \exists^{\geq 10} v.\,(\text{worksFor}(v,u) \wedge \text{employee}(v))$,
12. the contact person for a large scale project is a manager;
    $\forall u\forall v.\text{is_large}(u) \wedge \text{contactPerson}(u,v) \rightarrow \text{is_manager}(v)$.

Notice that Conjunct 2 contains an unrestricted occurrence of list$(u)$, thus program $\mathbb{LIST}$ is privileged in $\phi$.

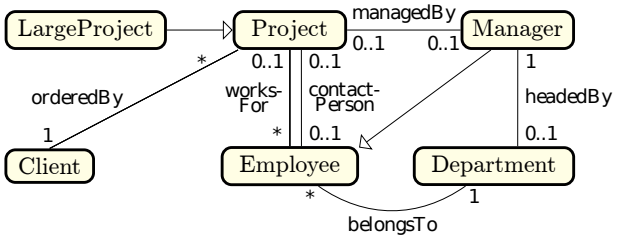


Fig. 4: A UML diagram for information system of a company as in [10].

# 4 Bounded model checking of pointer programs

Imperative pointer programs can naturally be viewed as state transition systems. A state stores data structures on heap and values of program variables in a given program location. Transitions correspond to program actions. A transition occurs between two states if the latter is obtained after successful execution of the corresponding action in the former state. In general the obtained transition system is infinite since it models program runs on every possible initial data structure on heap (e. g., a system representing list reversal program models its execution on every possible finite list). *Bounded Model Checking of Pointer Programs* [11] aims at discovering presence of `NULL`-pointer dereferences in pointer programs, but can also be used for violations of other safety properties. A counterexample is, roughly, a path from an initial state to a state where a `NULL`-pointer dereference occurs. We bound the length of paths we seek for, but the number of initial states remains unbounded. Program paths of bounded length are represented by universal two-variable formulas while admissible initial heap shapes are described by monadic Datalog programs. Satisfiability of the obtained formula of two-variable logic with Datalog is then equivalent to existence of a counterexample. Since logics we consider are decidable, so is the BMC of pointer programs. The present section is based on [12] (which is an extended version of [11]) with a modified presentation. The novelty lays in a generalization of the method: apart from checking for dangling pointers the BMC can now be used to discover variable aliasing, structure intersection or memory leaks.

## 4.1 Syntax of bounded pointer programs

A *bounded program* `BP` (also called a *straight-line program*) consists of two parts. The first one is a *struct declaration* specifying types of heap cells (called *templates* in [11]). Templates define pointers (*fields* in [11]) that start in a given heap cell. The second part is a finite sequence of actions specifying possible program executions.

A *struct declaration* is a finite directed graph with labeled edges. We call the vertices of this graph *types*, edge labels are called fields. Formally a struct declaration is a tuple $\langle T, f_1, f_2, \ldots, f_k \rangle$, where $T$ is a set of types and $f_1, f_2, \ldots, f_k$ are partial functions on $T$. Every allocated element of the heap has precisely one type $t \in T$. The meaning of $f_i(t_1) = t_2$ is that every heap element of type $t_1$ emits a pointer $f_i(\cdot, \cdot)$ to an element of type $t_2$ or to a special element `NULL`. Types $t$ are modeled by unary predicates $t(\cdot)$. Given a structure and a node $e$ of type $t$ in the structure we call $e$ a *$t$-cell*. Denote by struct(`BP`) the struct declaration of a bounded pointer program `BP`.

The set of *actions* $Act$ is defined by the grammar in Figure 5, where $t$ is a template, $s$ is a field, $x$ and $y$ are program variables, $e$ is a program variable or a constant `NULL`, and $\gamma$ is an arbitrary $\forall\forall$ formula that in particular may contain Boolean conditions over program variables and constants (including `NULL`) and the equality symbol $\approx$.

| | | |
|---|---|---|
| $Act ::=$ | $\texttt{assume}(\gamma)$ | *Skip to next action if condition $\gamma$ is satisfied, fail otherwise.* |
| $\mid$ | $y := e$ | *Assign the value $e$ to the variable $y$.* |
| $\mid$ | $y := s(x)$ | *Read the $s$-field of the cell pointed to by $x$ into $y$.* |
| $\mid$ | $s(x) := e$ | *Write $e$ to the $s$-field of the cell pointed to by $x$.* |
| $\mid$ | $\texttt{free}_t(x)$ | *Deallocate the $t$-cell pointed to by $x$.* |
| $\mid$ | $y := \texttt{new}_t()$ | *Allocate a new $t$-cell and assign its address to $y$.* |

Fig. 5: The action language.

Denote by actions(BP) the sequence of actions of a bounded pointer program BP. For an exemplary bounded pointer program and its struct declaration refer to Example 7.

**Semantics of bounded programs.** The semantics of actions and bounded programs is rather self explanatory and can be found in the thesis [43]. We will write $(\mathcal{A}, \alpha) \rightsquigarrow \mathcal{B}$ if $\mathcal{B}$ is obtained by executing an action $\alpha$ in state $\mathcal{A}$. We made only small changes compared to [11], like introduction of types that allows us to model (un)allocated elements of heap.

### 4.2 The Model Checking Problem for Bounded Programs

Given sets of pre-states and post-states specified by formulas of a logic with Datalog and a bounded program we want to check if an execution of the program in some pre-state leads to a post-state. This is formalized below using $\mathrm{C}^2$ + Datalog to specify pre- and post-states.

**Definition 4 (Model Checking for Bounded Programs).**
***Instance**: two* $\mathrm{C}^2$+Datalog *formulas $\varphi = [\mathbb{P}, \Delta]\phi$, $\varphi' = [\mathbb{P}', \Delta']\phi'$ and a bounded program* BP.
***Question**: does there exist a pre-state $\mathcal{A}$ and a post-state $\mathcal{B}'$ such that $\mathcal{A}_{\mathbb{P}} \models \varphi$, $(\mathcal{A}, \pi) \rightsquigarrow \mathcal{B}'$ and $\mathcal{B}'_{\mathbb{P}'} \models \varphi'$, where $\pi = \mathrm{actions}(\texttt{BP})$?*

In the above definition formula $\varphi$ is over some vocabularies $\Sigma_E$ and $\Sigma_I$, and $\varphi'$ is over fresh copies of these vocabularies, i. e., $\Sigma'_E$ and $\Sigma'_I$. Structures $\mathcal{A}$ and $\mathcal{B}$ are $\Sigma_E$-structures, and structure $\mathcal{B}'$ is obtained by renaming vocabulary of $\mathcal{B}$ to its primed version. We also assume that the vocabulary associated to BP is contained in $\Sigma_E$. If the question in the model checking problem for bounded programs has a positive answer then we say that the instance $\langle \varphi, \varphi', \texttt{BP} \rangle$ has a solution. In this section we assume that all dereference actions of BP of the form $y := s(x)$, $s(x) := e$ and $\texttt{free}_t(x)$ are prepended with allocation checks $\texttt{assume}(\mathit{alloc}(x))$, where $\mathit{alloc}(x)$ is a syntactic shortcut for the formula $\bigvee_{t \in T} t(x)$ denoting that $x$ is allocated. Note that these checks can be added automatically.

**Theorem 2 ([43], Theorem 4.4, Cor. 4.5).** *Model checking for bounded programs is polynomially reducible to finite satisfiability of* $\mathrm{C}^2$ + Datalog*, provided*

*that the total number of privileged Datalog programs in $\varphi$ and $\varphi'$ is at most 2. Therefore the problem is* NExpTime*-complete.*

The upper bound in the corollary above follows from the observation that the reduction is polynomial and the satisfiability for the specification logic is in NExpTime. The lower bound is obtained by a trivial reduction from satisfiability of the specification logic (take a formula $\varphi$ and create an instance $\langle \varphi, [\emptyset](x \approx x), x := x \rangle$).

### 4.3 Example analyses

Here we show that model checking for bounded programs can be applied to some common reasoning tasks, employed e. g., in optimizing compilers ([41]). One of such analyses is a question whether two pointer expressions may denote the same heap cell.

*Example 5 (Checking for variable aliasing).*
**Instance**: A $C^2$ + Datalog formula $\varphi$, a bounded program `BP` and program variables $x$, $y$.
**Question**: Does there exist a pre-state $\mathcal{A}$ such that $\mathcal{A}_{\mathbb{P}} \models \varphi$ and an execution of `BP` such that $x$ and $y$ reference the same heap cell in the post-state?
**Answer**: Reduction to model checking for bounded programs. Formula $\varphi$ and program `BP` are already defined. Define $\varphi' = [\emptyset](x' \approx y')$.

The question in the above example was about so called *may-aliasing*. Notice that in our setting we can also answer the *must-aliasing* question, i. e., if two variables refer to the same heap cell after each execution in each pre-state satisfying a formula; it is enough to ask if the model checking with the formula $x' \not\approx y'$ does not have solution. Although data structure traversals encoded by bounded pointer programs are deterministic, may- and must-aliasing are different problems: think of pre-states being a singly linked list with two non-`NULL` nodes $x$ and $y$. A bounded pointer program that moves $x$ and $y$ one element forward *may* produce a post-state where $x$ and $y$ are aliases, since it happens when these variables are already aliases in a pre-state. Clearly, they need not be aliases in every pre-state and the answer to must-aliasing problem is "no".

As we have mentioned `assume`$(alloc(x))$ can be employed to test for allocation of $x$ before dereferencing or deallocation. Since testing for correct dereferencing was one of the main problems solved in [11] we rephrase it as an instance of model checking for bounded programs. By $alloc'(u)$ we denote the syntactic shortcut for the primed version of $alloc(u)$, i. e., for the formula $\bigvee_{t \in T} t'(u)$.

*Example 6 (Checking for dereference of dangling or `NULL` pointers).*
**Instance**: A $C^2$ + Datalog formula $\varphi$, a bounded program `BP` and a program variable $x$.
**Question**: Does there exist a pre-state $\mathcal{A}$ such that $\mathcal{A}_{\mathbb{P}} \models \varphi$ and an execution of `BP` such that $x$ is a dangling or `NULL` pointer in the post-state?
**Answer**: Reduction to model checking for bounded programs. Formula $\varphi$ and program `BP` are already defined. Define $\varphi' = [\emptyset](\neg alloc'(x'))$.

The next example shows a more realistic extension of the example above. It is a slight modification of an example from [11]. A "real-life" pointer program gives rise to one or more (possibly infinitely many) bounded pointer programs obtained by choosing particular branches in conditional statements, by unwinding of loops and by inserting $\texttt{assume}(alloc(x))$ actions before dereferencing of $x$.

*Example 7 (Checking for dereference of a dangling or `NULL` node in a pointer program).* Figure 6 shows an example program $\texttt{PP}_{\texttt{cl}}$ taken from [11]. The struct declaration of the program is $\langle\{\texttt{cl_node}\}, \texttt{next}, \texttt{prev}\rangle$, where both `next` and `prev` are $\{\langle\texttt{cl_node}, \texttt{cl_node}\rangle\}$. Upon start it expects that the variable `c` points to a doubly linked circular list (realized by `next`- and `prev`-pointers). The program deallocates the cell pointed to by `c`, allocates a new cell and inserts it in place of the old one (using the temporary variables `nc` and `pc`).

```
{ clnode *nc; clnode *pc, clnode *c;
    nc:=next(c);
    pc:=prev(c);
    free_cl_node(c);
    c:=new_cl_node();
    next(c):=nc;
    prev(c):=pc;
    next(pc):=c;
    prev(nc):=c;
}
```

Fig. 6: Replacing an element in doubly-linked circular list.

The pointer program supplemented by allocation checks together with struct declaration for the program and formula $\varphi_{pre} = [\mathbb{P}](\phi \wedge \texttt{cl}(\texttt{c}))$ defining pre-states are presented in Figure 7. There are six bounded pointer programs of interest defined by $\texttt{PP}_{\texttt{cl}}$: $\texttt{BP}_{[0]-[1]}$, $\texttt{BP}_{[0]-[3]}$, $\texttt{BP}_{[0]-[6]}$, $\texttt{BP}_{[0]-[8]}$, $\texttt{BP}_{[0]-[10]}$ and $\texttt{BP}_{[0]-[12]}$, consisting of line ranges given in subscripts. For each of these bounded pointer programs we must check if the dereference occurring after the last line of the bounded program may fail due to dangling or `NULL` pointers. For example, to be sure that the dereference of $c$ in line 8 is correct we have to check that $c$ is allocated after the execution of $\texttt{BP}_{[0]-[6]}$. Thus we are interested in the the following six instances of the model checking problem for bounded programs.

$$\langle\varphi_{pre}, \texttt{BP}_{[0]-[1]}, [\emptyset](\neg alloc'(c'))\rangle, \langle\varphi_{pre}, \texttt{BP}_{[0]-[3]}, [\emptyset](\neg alloc'(c'))\rangle,$$
$$\langle\varphi_{pre}, \texttt{BP}_{[0]-[6]}, [\emptyset](\neg alloc'(c'))\rangle, \langle\varphi_{pre}, \texttt{BP}_{[0]-[8]}, [\emptyset](\neg alloc'(c'))\rangle,$$
$$\langle\varphi_{pre}, \texttt{BP}_{[0]-[10]}, [\emptyset](\neg alloc'(\texttt{pc}'))\rangle, \langle\varphi_{pre}, \texttt{BP}_{[0]-[12]}, [\emptyset](\neg alloc'(\texttt{nc}'))\rangle.$$

It turns out that $\texttt{PP}_{\texttt{cl}}$ is not pointer-safe: the instance with $\texttt{BP}_{[0]-[10]}$ has a solution. An analysis of the model of the corresponding formula reveals the reason. If `c` points to a circular list of length 1 then $\texttt{pc} \approx \texttt{c}$ after the second action, so `pc` is dangling after $\texttt{free}_{\texttt{cl}}(\texttt{c})$.

```
{ clnode *nc; clnode *pc,
  clnode *c;
 [0] assume(alloc(c));
 [1] nc:=next(c);
 [2] assume(alloc(c));
 [3] pc:=prev(c);
 [4] assume(alloc(c));
 [5] free_cl_node(c);
 [6] c:=new_cl_node();
 [7] assume(alloc(c));
 [8] next(c):=nc;
 [9] assume(alloc(c));
[10] prev(c):=pc;
[11] assume(alloc(pc));
[12] next(pc):=c;
[13] assume(alloc(nc));
[14] prev(nc):=c;
}
```

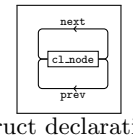


struct declaration

$$\begin{array}{l}\texttt{cl}(u) \leftarrow \texttt{cl_node}(u), u \approx \texttt{c}, \texttt{next}(u,v), \texttt{cl}'(v).\\ \texttt{cl}'(u) \leftarrow \texttt{cl_node}(u), u \approx \texttt{c}.\\ \texttt{cl}'(u) \leftarrow \texttt{cl_node}(u), u \not\approx \texttt{c}, \texttt{next}(u,v), \texttt{cl}'(v).\end{array}$$

Datalog program $\mathbb{P}$

$$\begin{array}{l}\forall u, v\big(\texttt{prev}(u,v) \Leftrightarrow \texttt{next}(v,u)\big) \wedge\\ \forall v \neg \texttt{next}(\texttt{NULL}, v) \wedge\\ \forall v \neg \texttt{prev}(\texttt{NULL}, v)\end{array}$$

axiom $\phi$

Fig. 7: Replacing an element in a doubly linked circular list by a new one; the initial condition that in pre-states c points to a doubly linked circular list is expressed by $[\mathbb{P}](\phi \wedge \texttt{cl}(\texttt{c}))$.

Problems from the examples above were already expressible using the logic from [11]. The logic was used to specify a pre-state; the specification of post-states was just a Boolean formula. By contrast, examples below employ more Datalog programs.

*Example 8 (Checking for structure intersection).*
**Instance**: A $C^2$ + Datalog formula $\varphi = [\mathbb{P}_1, \mathbb{P}_2]\phi$, where $\phi$ together with $\mathbb{P}_1$ (respectively $\mathbb{P}_2$) define some linked data structure by predicate $shape_1(\cdot)$ (respectively $shape_2(\cdot)$), at most one of $\mathbb{P}_1$ and $\mathbb{P}_2$ is privileged in $\varphi$, and a bounded program BP.
**Question**: Does there exist a pre-state $\mathcal{A}$ such that $\mathcal{A}_{\mathbb{P}_1,\mathbb{P}_2} \models \varphi$ and an execution of BP such that structures defined by $shape_1(\cdot)$ and $shape_2(\cdot)$ intersect in a non-NULL node in the post-state?
**Answer**: Reduction to model checking for bounded programs. Formula $\varphi$ and program BP are already defined. Let $\mathbb{P}'_1$, $\mathbb{P}'_2$ and $\phi'$ be $\mathbb{P}_1$ (respectively $\mathbb{P}_2$ and $\phi$) where all intensional and extensional are renamed to their primed versions (e. g., $shape_1(\cdot)$ becomes $shape'_1(\cdot)$). Define the formula $\varphi'$ as

$$\varphi' = [\mathbb{P}'_1, \mathbb{P}'_2](\phi' \wedge (\exists u\ shape'_1(u) \wedge shape'_2(u) \wedge u \not\approx \texttt{NULL}')).$$

Since our logic is closed under negation (it is enough to negate the first order part of a formula) we may also check for negation of the above properties, i. e., for non-dereference of dangling pointers, variable non-aliasing or structure non-intersection.

In the following example we employ a first-order interpretation of Datalog programs. A Datalog clause can be seen as a first-order implication. For a Datalog program $\mathbb{P}$ denote by $\overline{\mathbb{P}}$ the first order formula being the universally quantified conjunction of clauses in $\mathbb{P}$. Note that $\mathcal{M}_{\mathbb{P}}$ is a model of $\overline{\mathbb{P}}$, but the formula $\overline{\mathbb{P}}$ may also have other models, e. g., the structure $\mathcal{M}$ from Figure 1, whose all nodes are labeled by predicate list is a model of $\overline{\mathbb{P}_{\text{list}}}$, but is clearly distinct (i. e., greater) than $\mathcal{M}_{\mathbb{P}_{\text{list}}}$. We call $\overline{\mathbb{P}}$ the *the universal closure of* $\mathbb{P}$. By a simple transformation of formula $\overline{\mathbb{P}}$ one can obtain an equivalent $\text{FO}^2$ formula (see Proposition 2.25 in [43]). We will use this fact in Example 9 below.

We say that a bounded pointer program generates a *memory leak* if it creates a heap node which is allocated but unreachable from any of program variables.

*Example 9 (Checking for memory leaks).*
**Instance**: A $\text{C}^2$ + Datalog formula $\varphi_{pre} = [\mathbb{P}, \Delta]\phi_{pre}$ and a bounded program BP.
**Question**: Does there exist a pre-state $\mathcal{A}$, execution of BP and a node $a \in \mathcal{A}$ such that $\mathcal{A}_{\mathbb{P}} \models \varphi_{pre}$, the node $a$ is either unallocated or allocated and reachable from program variables in $\mathcal{A}$, and $a$ is allocated but unreachable in the post-state?
**Answer**: Reduction to model checking for bounded programs. Let $c_a$ be a fresh constant, which will be used to denote the above mentioned node $a \in \mathcal{A}$. Recall that the specification logic enforces the bounded-sharing restriction, which means, roughly, that only constant nodes may be shared by different pointers. Constant $c_a$ is an auxiliary symbol; formally it is a constant in the new extensional vocabulary but the node it interprets must not be shared unless it also interprets a constant from the old vocabulary $\Sigma_E$. To encode the above property we use a macro $bsr(u)$ defined as a conjunction of the following two formulas.

$$\bigwedge_{i\in\{1,\ldots,k\}} \bigwedge_{s_1(\cdot,\cdot)\in\Sigma(\mathbb{P}_i)} \forall u_1 \forall u_2 \; s_1(u_1,u) \wedge s_1(u_2,u) \wedge u_1 \neq u_2 \rightarrow \bigvee_{c\in\Sigma_E} u = c$$

$$\bigwedge_{i\in\{1,\ldots,k\}} \bigwedge_{s_1(\cdot,\cdot)\in\Sigma(\mathbb{P}_i)} \bigwedge_{s_2(\cdot,\cdot)\in\Sigma(\mathbb{P}_i)\setminus\{s_1\}} \forall u_1 \forall u_2 \; s_1(u_1,u) \wedge s_2(u_2,u) \rightarrow \bigvee_{c\in\Sigma_E} u = c$$

The instance of the model checking problem will be the tuple $\langle \varphi, \varphi', \text{BP}\rangle$, where $\varphi$ and $\varphi'$ are described below. Let $\text{struct}(\text{BP}) = \langle T, f_1, f_2, \ldots, f_k\rangle$ and $Var(\text{BP})$ be the set of program variables in BP.

Let $\text{reach}_{\text{pre}}(\cdot,\cdot)$ be a fresh intensional predicate and $\text{edge}(\cdot,\cdot)$ be a fresh extensional predicate. Define a Datalog program $\mathbb{Q}$ with two clauses $\text{reach}_{\text{pre}}(u) \leftarrow alloc(u) \wedge u \approx c_a$ and $\text{reach}_{\text{pre}}(u) \leftarrow alloc(u) \wedge edge(u,v) \wedge \text{reach}_{\text{pre}}(v)$ and a formula $\phi$ as $\forall u \forall v \; \text{edge}(u,v) \rightarrow \bigvee_{i=1}^{k} f_i(u,v)$. Intuitively, $\text{reach}_{\text{pre}}(u)$ means that $c_a$ is reachable from $u$. The following observation will be used to ensure that $c_a$ is reachable from a program variable $c$ in a pre-state: if $\mathcal{M}_{\mathbb{Q}}$ is any model of $[\mathbb{Q}](\phi \wedge \text{reach}_{\text{pre}}(c))$ then there is a path from $c$ to $c_a$ in $\mathcal{M}_{\mathbb{Q}}$ made of edges from $\{f_1(\cdot,\cdot), \ldots, f_k(\cdot,\cdot)\}$. We are now ready to define $\varphi$.

$$\varphi = [\mathbb{P}, \mathbb{Q}, \Delta]\left(\phi_{pre} \wedge \phi \wedge bsr(c_a) \wedge \left(\neg alloc(c_a) \vee \bigvee_{c\in Var(\text{BP})} reach_{pre}(c)\right)\right).$$

Let $\mathrm{reach}_{\mathrm{post}}(\cdot,\cdot)$ be a fresh intensional predicate. Define a Datalog program $\mathbb{R}$ with a clause $\mathrm{reach}_{\mathrm{post}}(u) \leftarrow \mathit{alloc}(u) \wedge u \approx c_a$ and clauses $\mathrm{reach}_{\mathrm{post}}(u) \leftarrow \mathit{alloc}(u) \wedge f_i(u,v) \wedge \mathrm{reach}_{\mathrm{post}}(v)$ for every $i \in \{1,\ldots,k\}$. The following observation will be used to ensure that $c_a$ is not reachable from a program variable $c$ in a post-state. Let $\overline{\mathbb{R}}$ be the $\mathrm{FO}^2$ formula equivalent to the universal closure of $\mathbb{R}$ (it exists by remark a the end of Section 3.1). If $\mathcal{M}$ is any model of $\overline{\mathbb{R}} \wedge \neg\mathrm{reach}_{\mathrm{post}}(c)$ then there is no path from $c$ to $c_a$ in $\mathcal{M}$ made of edges from $\{f_1(\cdot,\cdot),\ldots,f_k(\cdot,\cdot)\}$. In formula $\varphi'$ defined below we will use $\overline{\mathbb{R}}$ instead of $\mathbb{R}$ because the Datalog program $\mathbb{R}$ enforces bounded sharing on all edges $\{f_1(\cdot,\cdot),\ldots,f_k(\cdot,\cdot)\}$, while the first-order formula $\overline{\mathbb{R}}$ requires no such a restriction. This is important since formula $\varphi_{pre}$ describes pre-states by means of both Datalog program sequence $\mathbb{P}$ and the first order formula $\phi_{pre}$, and some of $\{f_1(\cdot,\cdot),\ldots,f_k(\cdot,\cdot)\}$ may appear only in $\phi_{pre}$ and therefore need not satisfy bounded-sharing for $\mathbb{P}$. We are now ready to define $\varphi'$.

$$\varphi' = [\emptyset]\left(\overline{\mathbb{R}'} \wedge \mathit{alloc}'(c'_a) \wedge \bigwedge_{c \in \mathit{Var}(\mathtt{BP})} \neg \mathit{reach}_{\mathit{post}}(c')\right).$$

In the formula above, which is just a $\mathrm{C}^2$ formula, $\mathbb{R}'$ is obtained from $\mathbb{R}$ by renaming all its extensional symbols $\{f_1(\cdot,\cdot),\ldots,f_k(\cdot,\cdot)\}$ to their primed versions $\{f'_1(\cdot,\cdot),\ldots,f'_k(\cdot,\cdot)\}$ and $c_a$ to $c'_a$, similarly for $\mathit{alloc}'(\cdot)$. The instance of the model checking problem is then $\langle \varphi, \varphi', \mathtt{BP}\rangle$.

We will now show that program BP generates memory leak when run on a state that satisfy $\varphi_{pre}$ if and only if the instance $\langle \varphi, \varphi', \mathtt{BP}\rangle$ has a solution. For the direct implication assume that BP runs on state $\mathcal{A}$, with $\varphi_{pre} \models \mathcal{A}$, and generates a memory leak. Therefore there exists $a \in \mathcal{A}$ such that either $a$ is unallocated or allocated and reachable from some constant node $c \in \mathcal{A}$ (recall that constant nodes of $\mathcal{A}$ model variables of program BP and NULL). Label node $a$ by a fresh constant $c_a$. If $a$ is allocated and reachable from $c$ then take an arbitrary path from $c$ to $a$ and label its edges by predicate $\mathrm{edge}(\cdot,\cdot)$. If $a$ is unallocated then we assign no $\mathrm{edge}(\cdot,\cdot)$ pointers. In both cases the obtained structure models the formula $\phi$. Let $\mathcal{A}_{\mathbb{Q}}$ be the least extension of the above modified $\mathcal{A}$ w. r. t. $\mathbb{Q}$. We will show that $\mathcal{A}_{\mathbb{Q}}$ satisfies $\varphi$. Clearly $\mathcal{A}_{\mathbb{Q}}$ satisfies $\phi_{pre}$ as, by assumption, $\mathcal{A} \models \phi_{pre}$. Similarly $\mathcal{A}_{\mathbb{Q}}$ satisfies $\phi$. Since $a$ is a node of $\mathcal{A}$ and $\mathcal{A}$ satisfies the bsr restriction, we also infer that $\mathcal{A}_{\mathbb{Q}}$ satisfies $\mathit{bsr}(c_a)$. If $a$ is unallocated then $\mathcal{A}_{\mathbb{Q}} \models \neg\mathit{alloc}(c_a)$. Otherwise, predicate $\mathrm{reach}_{\mathrm{pre}}(\cdot)$ labels a path from some constant $c$ to $c_a$. Therefore $\mathcal{A}_{\mathbb{Q}}$ satisfies $\varphi$. Let $\mathcal{B}'$ be a structure obtained after execution of BP on $\mathcal{A}$. Node $a$ is allocated, but unreachable from constants. Label $a$ by a fresh constant $c'_a$. Observe that $\mathcal{B}' \models \mathit{alloc}'(c'_a)$. Label each node of $\mathcal{B}'$ that is backward reachable from $c'_a$ by $\mathrm{reach}'_{\mathrm{post}}(\cdot)$. Then $\mathcal{B}' \models \overline{\mathbb{R}'}$. Since $a$ is reachable from no constant $c$, we have $\mathcal{B}' \models \neg\mathrm{reach}_{\mathrm{post}}(c')$ for all $c \in \mathit{Var}(\mathtt{BP})$. Therefore $\mathcal{B}' \models \varphi'$. Since $\mathcal{A}_{\mathbb{Q}} \models \varphi$ and $B' \models \varphi'$ the instance $\langle \varphi, \varphi', \mathtt{BP}\rangle$ has a solution.

Conversely, suppose that $\langle \varphi, \varphi', \mathtt{BP}\rangle$ has a solution with pre-state $\mathcal{A}$ and post-state $\mathcal{B}'$. Since $\mathcal{A} \models \varphi$, either the node $c_a$ is unallocated or some constant

$c \in \mathit{Var}(\mathtt{BP})$ is labeled with $\mathrm{reach}_{\mathrm{pre}}(\cdot)$. In the latter case there is a path from $c$ to $c_a$ labeled with $\mathrm{edge}(\cdot,\cdot)$, and then the assumption $\mathcal{A} \models \phi$ gives that the node $c_a$ is reachable from a variable. Therefore $c_a$ is either unallocated or reachable in the pre-state. Since $\mathcal{B}' \models \varphi'$, we have that $\mathcal{B}' \not\models \mathrm{reach}_{\mathrm{post}}(c')$ for all $c \in \mathit{Var}(\mathtt{BP})$. Observe that $\mathcal{B}'$ is some model of program $\mathbb{R}'$, so it contains the least model and thus it contains all atoms $\mathrm{reach}_{\mathrm{post}}(u)$ for all $u$ backward reachable from $c'_a$. Since it does not contain $\mathrm{reach}_{\mathrm{post}}(c')$, the node $c_a$ is not reachable from any program variable. But it is allocated and thus BP generates a memory leak.

## 5 Conclusions, open problems and future work

In this paper we extended the method of bounded model checking of pointer programs proposed in [11] by increasing the expressivity of the logic used for specification of data structures and properties of programs. We demonstrated expressivity of our logics on several examples. The examples provide an evidence of improvement over the method from [11] — it comes from extended expressibility of the underlying logics, which gives more sophisticated description of heaps (as in Examples 2, 3 and 4) and new analyses (Examples 8 and 9) not expressible in bounded model checking framework from [11]. Notice also that these analyses can be combined, provided that the number of privileged Datalog programs in the obtained instance of the model checking problem is at most 2.

Our method is based on translation to two-variable logic with counting quantifiers $\mathrm{C}^2$ with trees [14]. One may ask why we do not use directly this logic. The most important reason is that in Datalog it is relatively easy to express common data structures; the semantics based on least fixed points allows us to control in a simple way (a)cyclicity of these structures. Trying to express it directly in $\mathrm{C}^2$ with trees leads to formulas like our translations, which are too complicated to be used manually.

Relation of our logics with separation logics, $\mathrm{C}^2$ and $\mathrm{C}^2$ with trees raises a question about possibility of embedding decidable fragments of separation logics into these logics with counting quantifiers.

By using unary predicates and the $\mathtt{assume}(\gamma)$ construct we may model Boolean conditions in (finite unfoldings of) loops and conditional statements, provided that all data comes from a finite domain. We conjecture that a variant of the logic (the logic $\mathrm{C}^2 + \mathrm{Datalog}$ without privileged Datalog programs, which can be translated to $\mathrm{C}^2$ without trees) can be extended to a logic where data stored on heap can be accessed by equality tests and then translated to a decidable logic $\mathrm{C}^2$ with an equivalence relation [37]. This would allow us to extend the analyses expressible in $\mathrm{C}^2 + \mathrm{Datalog}$ to cope with data from infinite domains.

## References

1. T. Antonopoulos, N. Gorogiannis, C. Haase, M. I. Kanovich, and J. Ouaknine. Foundations for decision problems in separation logic with general inductive predicates. In A. Muscholl, editor, *Foundations of Software Science and Computation*

*Structures - 17th International Conference, FOSSACS 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, volume 8412 of *Lecture Notes in Computer Science*, pages 411–425. Springer, 2014.

2. G. Arnold, R. Manevich, M. Sagiv, and R. Shaham. Combining shape analyses by intersecting abstractions. In E. A. Emerson and K. S. Namjoshi, editors, *Verification, Model Checking, and Abstract Interpretation, 7th International Conference, VMCAI 2006, Charleston, SC, USA, January 8-10, 2006, Proceedings*, volume 3855 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 2006.
3. K. Bansal, R. Brochenin, and É. Lozes. Beyond shapes: Lists with ordered data. In L. de Alfaro, editor, *Foundations of Software Science and Computational Structures, 12th International Conference, FOSSACS 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*, volume 5504 of *Lecture Notes in Computer Science*, pages 425–439. Springer, 2009.
4. J. Berdine, C. Calcagno, and P. O'Hearn. A decidable fragment of separation logic. In *Proc. FSTTCS'04*, LNCS 3328, pages 97–109. Springer, 2004.
5. A. Biere, A. Cimatti, E. M. Clarke, and Y. Zhu. Symbolic model checking without bdds. In R. Cleaveland, editor, *TACAS*, volume 1579 of *Lecture Notes in Computer Science*, pages 193–207. Springer, 1999.
6. R. Brochenin, S. Demri, and É. Lozes. On the almighty wand. In M. Kaminski and S. Martini, editors, *Computer Science Logic, 22nd International Workshop, CSL 2008, 17th Annual Conference of the EACSL, Bertinoro, Italy, September 16-19, 2008. Proceedings*, volume 5213 of *Lecture Notes in Computer Science*, pages 323–338. Springer, 2008.
7. R. Brochenin, S. Demri, and É. Lozes. Reasoning about sequences of memory states. *Ann. Pure Appl. Logic*, 161(3):305–323, 2009.
8. J. Brotherston, C. Fuhs, J. A. N. Pérez, and N. Gorogiannis. A decision procedure for satisfiability in separation logic with inductive predicates. In Henzinger and Miller [20], page 25.
9. C. Calcagno, P. Gardner, and M. Hague. From separation logic to first-order logic. In V. Sassone, editor, *Foundations of Software Science and Computational Structures, 8th International Conference, FOSSACS 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings*, volume 3441 of *Lecture Notes in Computer Science*, pages 395–409. Springer, 2005.
10. D. Calvanese, T. Kotek, M. Simkus, H. Veith, and F. Zuleger. Shape and content: Incorporating domain knowledge into shape analysis. *CoRR*, abs/1312.6624, 2013.
11. W. Charatonik, L. Georgieva, and P. Maier. Bounded model checking of pointer programs. In *Proceedings of the 19th Annual Conference of the European Association for Computer Science Logic (CSL'05)*, pages 397–412, 2005.
12. W. Charatonik, L. Georgieva, and P. Maier. Bounded model checking of pointer programs. Technical Report MPI-I-2005-2-002, Max-Planck-Institut für Informatik, 2005.
13. W. Charatonik and P. Witkowski. On the complexity of the Bernays-Schönfinkel class with Datalog. In C. Fermüller and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 6397 of *Lecture Notes in Computer Science*, pages 187–201. Springer Berlin / Heidelberg, 2010.
14. W. Charatonik and P. Witkowski. Two-variable logic with counting and trees. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013,*

*New Orleans, LA, USA, June 25-28, 2013*, pages 73–82. IEEE Computer Society, 2013.
15. E. Clarke, D. Kroening, and F. Lerda. A tool for checking ANSI-C programs. In *Proc. TACAS'04*, LNCS 2988, pages 168–176. Springer, 2004.
16. B. Cook, C. Haase, J. Ouaknine, M. J. Parkinson, and J. Worrell. Tractable reasoning in a fragment of separation logic. In J. Katoen and B. König, editors, *CONCUR 2011 - Concurrency Theory - 22nd International Conference, CONCUR 2011, Aachen, Germany, September 6-9, 2011. Proceedings*, volume 6901 of *Lecture Notes in Computer Science*, pages 235–249. Springer, 2011.
17. P. Ferrara. Generic combination of heap and value analyses in abstract interpretation. In K. L. McMillan and X. Rival, editors, *Verification, Model Checking, and Abstract Interpretation - 15th International Conference, VMCAI 2014, San Diego, CA, USA, January 19-21, 2014, Proceedings*, volume 8318 of *Lecture Notes in Computer Science*, pages 302–321. Springer, 2014.
18. P. Ferrara, R. Fuchs, and U. Juhasz. TVAL+ : TVLA and value analyses together. In G. Eleftherakis, M. Hinchey, and M. Holcombe, editors, *Software Engineering and Formal Methods - 10th International Conference, SEFM 2012, Thessaloniki, Greece, October 1-5, 2012. Proceedings*, volume 7504 of *Lecture Notes in Computer Science*, pages 63–77. Springer, 2012.
19. P. Habermehl, R. Iosif, and T. Vojnar. Automata-based verification of programs with tree updates. In H. Hermanns and J. Palsberg, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 12th International Conference, TACAS 2006 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2006, Vienna, Austria, March 25 - April 2, 2006, Proceedings*, volume 3920 of *Lecture Notes in Computer Science*, pages 350–364. Springer, 2006.
20. T. A. Henzinger and D. Miller, editors. *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*. ACM, 2014.
21. M. Huth and S. Pradhan. Consistent partial model checking. *Electronic Notes in Theoretical Computer Science*, 23, 2003.
22. N. Immerman, A. Rabinovich, T. Reps, M. Sagiv, and G. Yorsh. Verification via structure simulation. In *Proc. CAV'04*, LNCS 3114, pages 281–294. Springer, 2004.
23. R. Iosif, A. Rogalewicz, and J. Simácek. The tree width of separation logic with recursive definitions. In M. P. Bonacina, editor, *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction, Lake Placid, NY, USA, June 9-14, 2013. Proceedings*, volume 7898 of *Lecture Notes in Computer Science*, pages 21–38. Springer, 2013.
24. D. Jackson and M. Vaziri. Finding bugs with a constraint solver. In *Proc. ISSTA'00*, pages 14–25, 2000.
25. N. Klarlund and M. I. Schwartzbach. Graph types. In *Proc. POPL'93*, pages 196–205, 1993.
26. T. Kotek, M. Simkus, H. Veith, and F. Zuleger. Towards a description logic for program analysis: Extending ALCQIO with reachability. In M. Bienvenu, M. Ortiz, R. Rosati, and M. Simkus, editors, *Informal Proceedings of the 27th International Workshop on Description Logics, Vienna, Austria, July 17-20, 2014.*, volume 1193 of *CEUR Workshop Proceedings*, pages 591–594. CEUR-WS.org, 2014.
27. V. Kuncak, P. Lam, and M. C. Rinard. Role analysis. In J. Launchbury and J. C. Mitchell, editors, *Conference Record of POPL 2002: The 29th SIGPLAN-SIGACT*

*Symposium on Principles of Programming Languages, Portland, OR, USA, January 16-18, 2002*, pages 17–32. ACM, 2002.

28. V. Kuncak and M. C. Rinard. Generalized records and spatial conjunction in role logic. In R. Giacobazzi, editor, *Static Analysis, 11th International Symposium, SAS 2004, Verona, Italy, August 26-28, 2004, Proceedings*, volume 3148 of *Lecture Notes in Computer Science*, pages 361–376. Springer, 2004.
29. S. K. Lahiri and S. Qadeer. Back to the future: revisiting precise program verification using SMT solvers. In G. C. Necula and P. Wadler, editors, *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, pages 171–182. ACM, 2008.
30. K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In E. M. Clarke and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning - 16th International Conference, LPAR-16, Dakar, Senegal, April 25-May 1, 2010, Revised Selected Papers*, volume 6355 of *Lecture Notes in Computer Science*, pages 348–370. Springer, 2010.
31. R. Manevich, E. Yahav, G. Ramalingam, and S. Sagiv. Predicate abstraction and canonical abstraction for singly-linked lists. In R. Cousot, editor, *Verification, Model Checking, and Abstract Interpretation, 6th International Conference, VMCAI 2005, Paris, France, January 17-19, 2005, Proceedings*, volume 3385 of *Lecture Notes in Computer Science*, pages 181–198. Springer, 2005.
32. A. Møller and M. I. Schwartzbach. The pointer assertion logic engine. In *Proc. PLDI'01*, pages 221–231, 2001.
33. H. H. Nguyen, C. David, S. Qin, and W. Chin. Automated verification of shape and size properties via separation logic. In B. Cook and A. Podelski, editors, *Verification, Model Checking, and Abstract Interpretation, 8th International Conference, VMCAI 2007, Nice, France, January 14-16, 2007, Proceedings*, volume 4349 of *Lecture Notes in Computer Science*, pages 251–266. Springer, 2007.
34. P. W. O'Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In L. Fribourg, editor, *Computer Science Logic, 15th International Workshop, CSL 2001. 10th Annual Conference of the EACSL, Paris, France, September 10-13, 2001, Proceedings*, volume 2142 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2001.
35. E. Pek, X. Qiu, and P. Madhusudan. Natural proofs for data structure manipulation in C using separation logic. In M. F. P. O'Boyle and K. Pingali, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14, Edinburgh, United Kingdom - June 09 - 11, 2014*, page 46. ACM, 2014.
36. R. Piskac, T. Wies, and D. Zufferey. Automating separation logic with trees and data. In A. Biere and R. Bloem, editors, *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings*, volume 8559 of *Lecture Notes in Computer Science*, pages 711–728. Springer, 2014.
37. I. Pratt-Hartmann. Logics with counting and equivalence. In Henzinger and Miller [20], page 76.
38. X. Qiu, P. Garg, A. Stefanescu, and P. Madhusudan. Natural proofs for structure, data, and separation. In H. Boehm and C. Flanagan, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*, pages 231–242. ACM, 2013.

39. Z. Rakamaric, R. Bruttomesso, A. J. Hu, and A. Cimatti. Verifying heap-manipulating programs in an SMT framework. In K. S. Namjoshi, T. Yoneda, T. Higashino, and Y. Okamura, editors, *Automated Technology for Verification and Analysis, 5th International Symposium, ATVA 2007, Tokyo, Japan, October 22-25, 2007, Proceedings*, volume 4762 of *Lecture Notes in Computer Science*, pages 237–252. Springer, 2007.
40. A. Rensink. Canonical graph shapes. In *Proc. ESOP'04*, LNCS 2986, pages 401–415. Springer, 2004.
41. T. Reps, M. Sagiv, and R. Wilhelm. Shape analysis and applications. In Y. N. Srikant and P. Shankar, editors, *The Compiler Design Handbook: Optimizations and Machine Code Generation, Second Edition*. CRC Press, Inc., Boca Raton, FL, USA, 2nd edition, 2007.
42. M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape-analysis problems via 3-valued logic. *ACM TOPLAS*, 24(2):217–298, 2002.
43. P. Witkowski. *Complexity of Some Logics Extended with Monadic Datalog Programs*. PhD thesis, Institute of Computer Science, University of Wrocław, 2014. `http://www.ii.uni.wroc.pl/~pwit/thesis/thesis.pdf`.
44. G. Yorsh, A. Rabinovich, M. Sagiv, A. Meyer, and A. Bouajjani. A logic of reachable patterns in linked data-structures. *Journal of Logic and Algebraic Programming*, 73(1-2):111 – 142, 2007. Foundations of Software Science and Computation Structures 2006 (FOSSACS 2006).
45. G. Yorsh, T. Reps, and M. Sagiv. Symbolically computing most-precise abstract operations for shape analysis. In *Proc. TACAS'04*, LNCS 2988, pages 530–545. Springer, 2004.
46. K. Zee, V. Kuncak, and M. C. Rinard. Full functional verification of linked data structures. In R. Gupta and S. P. Amarasinghe, editors, *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation, Tucson, AZ, USA, June 7-13, 2008*, pages 349–361. ACM, 2008.
47. K. Zee, V. Kuncak, and M. C. Rinard. An integrated proof language for imperative programs. In M. Hind and A. Diwan, editors, *Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2009, Dublin, Ireland, June 15-21, 2009*, pages 338–351. ACM, 2009.